arXiv:1604.00617v1 [cs.NA] 3 Apr 2016

A Direct Elliptic Solver Based on Hierarchically Low-rank Schur Complements

Gustavo Chávez, George Turkiyyah, and David Keyes

1 Introduction

Cyclic reduction was conceived for the solution of tridiagonal linear systems, such as the one-dimensional Poisson equation [11]. Generalized to higher dimensions, it is known as block cyclic reduction (BCR) [4]. It can be used for general (block) Toeplitz and (block) tridiagonal linear systems; however, it is not competitive for large problems, because its arithmetic complexity grows superlinearly. Cyclic reduction can be thought of as a direct Gaussian elimination that recursively computes the Schur complement of half of the system. Schur complement computations have complexity that grows with the cost of the inverse, but by considering a tridiagonal system and an even/odd ordering, cyclic reduction can decouple the system in such a way that the inverse of a large block is the block-wise inverse of a collection of independent smaller blocks. This addresses the most expensive step of the Schur complement computation in terms of operation complexity and does so in a way that launches concurrent subproblems. Its concurrency feature in the form of recursive bisection makes it interesting for parallel environments, provided that its arithmetic complexity can be improved.

We address the time and memory complexity growth of the traditional cyclic reduction algorithm by approximating dense blocks as they arise with hierarchical matrices ($\mathcal{H}$-matrices). The effectiveness of the block approximation relies on the rank structure of the original matrix. Many relevant operators are known to have low rank off the diagonal. This philosophy follows recent work discussed below, but to our knowledge this is the first demonstration of the utility of complexity-reducing hierarchical substitution in the context of cyclic reduction.

{gustavo.chavezchavez,george.turkiyyah,david.keyes}@kaust.edu.sa, Extreme Computing Research Center, King Abdullah University of Science and Technology, Thuwal 23955, Saudi Arabia. (Submitted to the Proceedings of the 23th International Conference on Domain Decomposition Methods in Jeju, Korea)

The synergy of cyclic reduction and hierarchical matrices leads to a parallel fast direct solver of log-linear arithmetic complexity, $O(N \log^2 N)$, with controllable accuracy, using cyclic reduction outside and $\mathcal{H}$-Matrix arithmetic operations inside. The algorithm is purely algebraic, depending only on a block tridiagonal structure. Neither geometric information nor sophisticated re-orderings of the matrix are required. We call it Accelerated Cyclic Reduction (ACR). Using a well-known implementation of $\mathcal{H}$-LU, we demonstrate the range of applicability of ACR over a set of model problems including strong convection and high frequency Helmholtz, including variable coefficient problems that cannot be tackled with the traditional FFT enabled version of cyclic reduction, FACR [19]. We show that it is competitive in time to solution with a global $\mathcal{H}$-LU factorization that does not exploit the cyclic reduction structure. The fact that ACR is completely algebraic expands its range of applicability to problems with arbitrary coefficient structure within the block tridiagonal sparsity structure, subject to their amenability to rank compression. This gives the method robustness in some applications that are difficult for multigrid. The concurrency and flexibility to tune the accuracy of individual matrix block approximations makes it interesting for emerging manycore architectures. Finally, as with many direct solvers, there are complexity-accuracy tradeoffs that lead to the development of a new scalable preconditioner.

2 Related Work

Exploiting underlying low-rank structure is a trending strategy for improving the performance of sparse direct solvers.

Nested dissection based clustering of an $\mathcal{H}$-matrix is known as $\mathcal{H}$-Cholesky by Ibragimov *et al.* [12] and $\mathcal{H}$-LU by Grasedyck *et al.* [7, 8], the main idea being to introduce $\mathcal{H}$-Matrix approximation on Schur complements based on domain decomposition. This is accomplished by a nested dissection ordering of the unknowns, and the advantage is that large blocks of zeros are preserved after factorization. The non-zero blocks are replaced with low-rank approximations, and an LU factorization is performed, substituting in $\mathcal{H}$ operations. Recently, Kriemann *et al.* [13] demonstrated that $\mathcal{H}$-LU implemented with a task-based scheduling based on a directed acyclic graph is well suited for modern many-core systems when compared with the conventional recursive algorithm. A similar line of work by Xia *et al.* [22] also proposes the construction of a rank-structured Cholesky factorization via the HSS hierarchical format. Figure 1 illustrates the differences between nested dissection ordering and the even/odd (or red/black) ordering of cyclic reduction.

Multifrontal factorization, with low-rank approximations of frontal matrices, as in the work of Xia *et al.* [20] also relies on nested dissection as the permutation strategy, but it uses the multifrontal method as a solver.

Frontal matrices are approximated with the HSS format, while the solver relies on the corresponding HSS algorithms for elimination [21]. A similar line of work is the generalization of this method to 3D problems and general meshes by Schmitz *et al.* [18, 17]. More recently, Ghysels *et al.* [6] introduced a method based on a fast ULV decomposition [5] and randomized sampling of HSS matrices in a many-core environment, where HSS approximations are used to approximate fronts of large enough size, as the complexity constant in building an HSS approximation is only convenient for large matrices.

This strategy is not limited to any specific hierarchical format. Aminfar *et al.* [3] proposed the use of the HODLR matrix format [1], also in the context of the multifrontal method. The well known solver MUMPS now also exploits the low-rank property of frontal matrices to accelerate its multifrontal implementation, as described in [2].

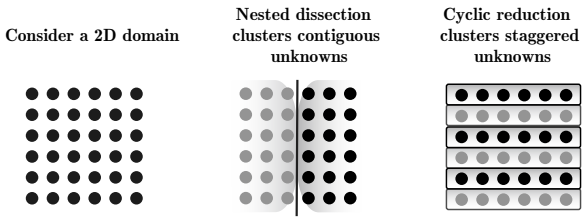

Fig. 1 The nested dissection ordering recursively clusters contiguous unknowns by bisection, whereas the red/black ordering recursively clusters staggered unknowns, allowing isolation of a new readily manipulated diagonal block.

3 Accelerated Cyclic Reduction

Consider the two-dimensional linear variable-coefficient Poisson equation (1) and its corresponding block tridiagonal matrix structure resulting from a second order finite difference discretization, as shown in (2):

$$-\nabla \cdot \kappa(\mathbf{x}) \nabla u = f(\mathbf{x}), \tag{1}$$

$$A = \text{tridiag}(E_i, D_i, F_i) = \begin{bmatrix} D_1 & F_1 & & & \\ E_2 & D_2 & F_2 & & \\ & \ddots & \ddots & \ddots & \\ & & E_{n-1} & D_{n-1} & F_{n-1} \\ & & & E_n & D_n \end{bmatrix}. \tag{2}$$

We leverage the fact that for arbitrary $\kappa(\mathbf{x})$, the blocks D_i are *exactly* representable by rank 1 $\mathcal{H}$-Matrices, and the blocks E_i and F_i are diagonal. As

cyclic reduction progresses, the resulting blocks will have a bounded increase in the numerical ranks of their off-diagonal blocks. This numerical off-diagonal rank may be tuned to accommodate for a specified accuracy. We choose the $\mathcal{H}$-Matrix format proposed in [9] by Hackbusch, although ACR is not limited to a specific hierarchical format. In terms of admissibility condition, we choose weak admissibility, as the sparsity structure is known beforehand and it proved effective in our numerical experiments.

Approximating each block as an $\mathcal{H}$-matrix, we use the corresponding hierarchical arithmetic operations as cyclic reduction progresses, instead of the conventional linear algebra arithmetic operations. The following table summarizes the complexity estimates in terms of time and memory while dealing with a $n \times n$ block in a typical dense format and as a block-wise approximation with a rank-r $\mathcal{H}$-matrix.

	Inverse	Storage
Dense Block	$\mathcal{O}(n^3)$	$\mathcal{O}(n^2)$
$\mathcal{H}$ Block	$\mathcal{O}(r^2 n \log^2 n)$	$\mathcal{O}(rn \log n)$

The following table summarizes the complexity estimates of the methods discussed so far in two dimensions, neglecting the dependence upon rank.

	Operations	Memory
BCR	$\mathcal{O}(N^2)$	$\mathcal{O}(N^{1.5} \log N)$
$\mathcal{H}$-LU	$\mathcal{O}(N \log^2 N)$	$\mathcal{O}(N \log N)$
ACR	$\mathcal{O}(N \log^2 N)$	$\mathcal{O}(N \log N)$

With block-wise approximations in place, block cyclic reduction becomes ACR. BCR consists of two phases: reduction and back-substitution. The reduction phase is equivalent to block Gaussian elimination without pivoting on a permuted system $(PAP^T)(Pu) = Pf$. Permutation decouples the system, and the computation of the Schur complement reduces the problem size by half. This process is recursive and finishes when a single block is reached, although the recursion can be stopped when the system is small enough to be solved directly. The second phase performs a conventional back-substitution to find the solution at every point of the domain.

As an illustration, consider a system of $n = 8$ points per dimension, which translates into a $N \times N$ sparse matrix, with $N = n^2$. The first step is to permute the system, which with an even/odd ordering becomes:

$$\left[\begin{array}{cccc|cccc} D_0 & & & & F_0 & & & \\ & D_2 & & & E_2 & F_2 & & \\ & & D_4 & & & E_4 & F_4 & \\ & & & D_6 & & & E_6 & F_6 \\ \hline E_1 & F_1 & & & D_1 & & & \\ & E_3 & F_3 & & & D_3 & & \\ & & E_5 & F_5 & & & D_5 & \\ & & & E_7 & & & & D_7 \end{array}\right] \left[\begin{array}{c} u_0 \\ u_2 \\ u_4 \\ u_6 \\ \hline u_1 \\ u_3 \\ u_5 \\ u_7 \end{array}\right] = \left[\begin{array}{c} f_0 \\ f_2 \\ f_4 \\ f_6 \\ \hline f_1 \\ f_3 \\ f_5 \\ f_7 \end{array}\right]. \tag{3}$$

Consider the above 2×2 partitioned system as H. The upper-left block is block-diagonal, which means that its inverse can be computed as the inverse of each individual block (D_0, D_2, D_4, and D_6), in parallel and with hierarchical matrix arithmetics. The Schur complement of the upper-left partition may then be computed as follows:

$$\left[\begin{array}{c|c} H_{11} & H_{12} \\ \hline H_{21} & H_{22} \end{array}\right] \left[\begin{array}{c} u_{even} \\ \hline u_{odd} \end{array}\right] = \left[\begin{array}{c} f_{even} \\ \hline f_{odd} \end{array}\right]. \tag{4}$$

$$(H_{22} - H_{21}H_{11}^{-1}H_{12})u_{odd} = f^{(1)}, \qquad f^{(1)} = f_{odd} - H_{21}H_{11}^{-1}f_{even}. \tag{5}$$

Superscripts indicates algorithmic steps. A key property of the Schur complement of a block tridiagonal matrix is that it also yields a block tridiagonal matrix, as can been seen in the resulting permuted matrix system:

$$\left[\begin{array}{ccc|ccc} D_0^{(1)} & & & F_0^{(1)} & & \\ & D_2^{(1)} & & E_2^{(1)} & F_2^{(1)} & \\ \hline E_1^{(1)} & F_1^{(1)} & & D_1^{(1)} & & \\ & E_3^{(1)} & & & D_3^{(1)} & \end{array}\right] \left[\begin{array}{c} u_0^{(1)} \\ u_2^{(1)} \\ \hline u_1^{(1)} \\ u_3^{(1)} \end{array}\right] = \left[\begin{array}{c} f_0^{(1)} \\ f_2^{(1)} \\ \hline f_1^{(1)} \\ f_3^{(1)} \end{array}\right]. \tag{6}$$

One step further, the computation of the Schur complement results in:

$$\left[\begin{array}{cc} D_0^{(2)} & F_0^{(2)} \\ E_1^{(2)} & D_1^{(2)} \end{array}\right] \left[\begin{array}{c} u_0^{(2)} \\ u_1^{(2)} \end{array}\right] = \left[\begin{array}{c} f_0^{(2)} \\ f_1^{(2)} \end{array}\right]. \tag{7}$$

A last round of permutation and Schur complement computation leads to the $D_0^{(3)}$ block, which is the last step of the reduction phase of Cyclic Reduction. A conventional back-substitution step recovers the solution.

4 Numerical Results

We select two test cases to provide a baseline of performance and robustness as compared with high-performance implementations of $\mathcal{H}$-LU and AMG, namely HLIBpro [10], and Hypre [14]. Tests are performed in the shared memory environment of a 12-core Intel Westemere processor.

The first test is the wave Helmholtz equation. For large values of kh, where h is the mesh spacing, discretization leads to an indefinite matrix. Performance over a range of k is shown in Figure 2, for $h = 2^{-10}$.

$$\begin{aligned} \nabla^2 u + k^2 u = f(\mathbf{x}), \qquad & \mathbf{x} \in \Omega = [0,1]^2 \quad u(\mathbf{x}) = 0,\ x \in \Gamma \\ & f(\mathbf{x}) = 100 e^{-100((x-0.5)^2+(y-0.5)^2)}. \end{aligned} \tag{8}$$

The second test problem is the convection-diffusion equation with recirculating flow. The discretization of this equation leads to a nonsymmetric matrix.

$$\begin{aligned} -\nabla^2 u + \alpha b(\mathbf{x}) \cdot \nabla u = f(\mathbf{x}), \qquad & \mathbf{x} \in \Omega = [0,1]^2 \;\; u(\mathbf{x}) = 0, \; x \in \Gamma \\ b = \begin{pmatrix} \cos(8\pi x) \\ \sin(8\pi y) \end{pmatrix} \qquad & f(\mathbf{x}) = 100 e^{-100((x-0.5)^2+(y-0.5)^2)}. \end{aligned} \tag{9}$$

We progressively increase the convection dominance with α, accentuating the skew-symmetry. For small α AMG out performs ACR, but since the increase in α does not depend on the rank structure of the matrix, ACR maintains its performance for any α. Performance can be seen in Figure 3

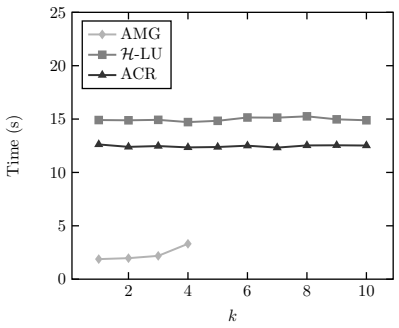

Fig. 2 Solvers performance while decreasing the number of points per wavelength. AMG fails to converge for large k.

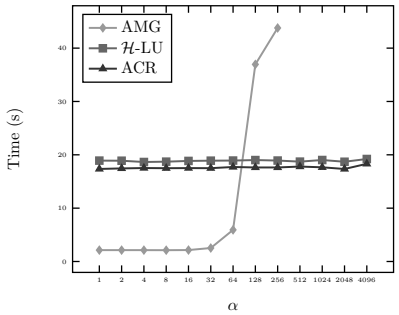

Fig. 3 Stability of the solution as the convection dominance increases. AMG fails to converge for large α.

The discretization of 3D elliptic operators also leads to a block tridiagonal structure, with the difference that each block is of size $n^2 \times n^2$, as opposed to $n \times n$ as in the 2D discretization. Also, the numerical rank of the off-diagonal blocks grows faster than in the 2D case, which would lead into a superlinear solver of $\mathcal{O}(N^{\frac{4}{3}} \log N)$. This complexity is not practical for 3D problems. Nonetheless, it is possible to use CR as a preconditioner, as proposed in [16, 15]. It is likewise advantageous to use ACR as a preconditioner since it is possible to control the accuracy of the factorization, which directly translates into a tunable parameter of performance, while preserving the rich concurrency features of the method. Preconditioning represents an interesting extension of the usability of ACR to tackle three dimensional problems, with low-rank parameterization techniques. Further details are discussed in a companion paper in preparation with a larger set of test cases in two and three dimensions.

5 Concluding Remarks

We present a fast direct solver, ACR, for structured sparse linear systems that arise from the discretization of 2D elliptic operators. The solver approximates every block using an $\mathcal{H}$-Matrix, resulting in a log-linear arithmetic complexity of $\mathcal{O}(N \log^2 N)$ with memory requirements of $\mathcal{O}(N \log N)$.

Robustness and applicability are demonstrated on model scalar problems and contrasted with established solvers based on the $\mathcal{H}$-LU factorization and algebraic multigrid. Multigrid maintains superiority in scalar problems with sufficient definiteness and symmetry, whereas hierarchical matrix based solvers can tackle problems where these properties are lacking.

Although being of the same asymptotic complexity as $\mathcal{H}$-LU, ACR has fundamentally different algorithmic roots which produce a novel alternative for a relevant class of problems with competitive performance, increasing concurrency as the problem grows, and almost optimal memory requirements. In a companion paper we re-introduce the consideration of cyclic reduction as a preconditioner by exploiting tunable accuracy of low-rank approximation and the exploitation of multicore features.

References

1. S. Ambikasaran and E. Darve. An $\mathcal{O}(N \log N)$ fast direct solver for partial hierarchically semiseparable matrices. *Journal of Scientific Computing*, 57(3):477–501, Dec 2013.
2. P. Amestoy, A. Buttari, G. Joslin, J.-Y. L'Excellent, M. Sid-Lakhdar, C. Weisbecker, M. Forzan, C. Pozza, R. Perrin, and V. Pellissier. Shared-memory parallelism and low-rank approximation techniques applied to direct solvers in FEM simulation. *IEEE Transactions on Magnetics*, 50(2):517–520, Feb 2014.
3. A. Aminfar and E. Darve. A fast sparse solver for Finite-Element matrices. *arXiv:1403.5337 [cs.NA]*, pages 1–25, 2014.
4. B. L. Buzbee, G. H. Golub, and C. W. Nielson. On direct methods for solving Poisson equation. *SIAM Journal on Numerical Analysis*, 7(4):pp. 627–656, 1970.
5. S. Chandrasekaran, M. Gu, and T. Pals. A fast ULV decomposition solver for hierarchically semiseparable representations. *SIAM J. Matrix Anal. Appl.*, 28(3):603–622, Aug 2006.
6. P. Ghysels, X. S. Li, F.-H. Rouet, S. Williams, and A. Napov. An efficient multicore implementation of a novel HSS-structured multifrontal solver using randomized sampling. *arXiv:1502.07405 [cs.MS]*, pages 1–26, 2015.
7. L. Grasedyck, R. Kriemann, and S. Le Borne. Parallel black box $\mathcal{H}$-LU preconditioning for elliptic Boundary Value Problems. *Computing and Visualization in Science*, 11(4-6):273–291, 2008.
8. L. Grasedyck, R. Kriemann, and S. Le Borne. Domain decomposition based $\mathcal{H}$-LU preconditioning. *Numerische Mathematik*, 112(4):565–600, 2009.
9. W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-Matrices. Part I: Introduction to $\mathcal{H}$-Matrices. *Computing*, 62(2):89–108, 1999.
10. Wolfgang Hackbusch, Steffen Börm, and Lars Grasedyck. HLib 1.4. `http://hlib.org`, 1999-2012. Max-Planck-Institut, Leipzig.

11. R. W. Hockney. A fast direct solution of Poisson's equation using Fourier analysis. *J. ACM*, 12(1):95–113, Jan 1965.
12. I. Ibragimov, S. Rjasanow, and K. Straube. Hierarchical Cholesky decomposition of sparse matrices arising from curl-curl-equation. *Journal of Numerical Mathematics*, 15(1):31–57, 2007.
13. R. Kriemann. $\mathcal{H}$-LU factorization on many-core systems. *Computing and Visualization in Science*, 16(3):105–117, 2013.
14. Lawrence Livermore National Laboratory. *hypre: High Performance Preconditioners*. `http://www.llnl.gov/CASC/hypre/`.
15. A. Reusken. On the approximate cyclic reduction preconditioner. *SIAM J. Sci. Comput*, 21:565–590, 2000.
16. G. Rodrigue and D. Wolitzer. Preconditioning by incomplete block cyclic reduction. *Mathematics of Computation*, 42(166):549–565, 1984.
17. P. G. Schmitz and L. Ying. A fast direct solver for elliptic problems on general meshes in 2D. *Journal of Computational Physics*, 231(4):1314–1338, 2012.
18. P. G. Schmitz and L. Ying. A fast nested dissection solver for Cartesian 3D elliptic problems using hierarchical matrices. *Journal of Computational Physics*, 258:227–245, 2014.
19. P. Swarztrauber. The methods of Cyclic reduction, Fourier analysis and the FACR algorithm for the discrete solution of Poisson equation on a rectangle. *SIAM Review*, 19(3):490–501, 1977.
20. J. Xia, S. Chandrasekaran, M. Gu, and X. Li. Superfast multifrontal method for large structured linear systems of equations. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1382–1411, 2010.
21. J. Xia, S. Chandrasekaran, M. Gu, and X. S. Li. Fast algorithms for hierarchically semiseparable matrices. *Numerical Linear Algebra with Applications*, 17(6):953–976, 2010.
22. J. Xia and M. Gu. Robust approximate Cholesky factorization of rank-structured symmetric positive definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 31(5):2899–2920, 2010.